

Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) **EP 0 872 808 A1**

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:  
21.10.1998 Bulletin 1998/43

(51) Int. Cl.<sup>6</sup>: **G06K 11/08**

(21) Application number: 98106380.3

(22) Date of filing: 07.04.1998

(84) Designated Contracting States:  
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE**  
Designated Extension States:  
**AL LT LV MK RO SI**

(30) Priority: 18.04.1997 US 844444


(71) Applicant:  
**DIGITAL EQUIPMENT CORPORATION  
Maynard Massachusetts 01754-1418 (US)**

(72) Inventors:  
• **Waters, Keith**  
**West Newton, Massachusetts 02165 (US)**  
• **Rehg, James M.**  
**Arlington, Massachusetts 02174 (US)**  
• **Loughlin, Maria**  
**Belmont, Massachusetts 02178 (US)**  
• **Kang, Sing Bing**  
**Cambridge, Massachusetts 02138 (US)**

(74) Representative: **Betten & Resch**  
**Reichenbachstrasse 19**  
**80469 München (DE)**

(54) **Method and apparatus for visual sensing of humans for active public interfaces**

(57) An active public user interface in a computerized kiosk senses humans visually using movement and color to detect changes in the environment indicating the presence of people. Interaction spaces are defined and the system records an initial model of its environment which is updated over time to reflect the addition or subtraction of inanimate objects and to compensate for lighting changes. The system develops models of the moving objects and is thereby able to track people as they move about the interaction spaces. A stereo camera system further enhances the system's ability to sense location and movement. The kiosk presents audio and visual feedback in response to what it "sees."



**FIG. 2**

**EP 0 872 808 A1**

## Description

### FIELD OF THE INVENTION

This invention relates generally to computer systems, and more particularly to computerized human-computer interfaces.

### BACKGROUND OF THE INVENTION

Computer vision-based sensing of users enables a new class of public multi-user computer interfaces. An interface such as an automated information dispensing kiosk represents a computing paradigm that differs from the conventional desktop environment and correspondingly requires a user interface that is unlike the traditional Window, Icon, Mouse and Pointer (WIMP) interface. Consequently, as user interfaces evolve and migrate off the desktop, vision-based human sensing will play an increasingly important role in human-computer interaction.

Human sensing techniques that use computer vision can play a significant role in public user interfaces for kiosk-like computerized appliances. Computer vision using unobtrusive video cameras can provide a wealth of information about users, ranging from their three dimensional location to their facial expressions, and body posture and movements. Although vision-based human sensing has received increasing attention, relatively little work has been done on integrating this technology into functioning user interfaces.

The dynamic, unconstrained nature of a public space, such as a shopping mall, poses a challenging user interface problem for a computerized kiosk. This user interface problem can be referred to as the public user interface problem, to differentiate it from interactions that take place in a structured, single-user desktop environments. A fully automated public kiosk interface must be capable of actively initiating and terminating interactions with users. The kiosk must also be capable of dividing its resources among multiple users in an equitable manner.

The prior art technique for sensing users as applied in the Alive system is described in "Pfunder: Real-time Tracking of the Human Body," Christopher Wren, Ali Azarbayejani, Trevor Darrell, and Alex Pentland, IEEE 1996. Another prior art system is described in "Real-time Self-calibrating Stereo Person Tracking Using 3-D Shape Estimation from Blob Features," Ali Azarbayejani and Alex Pentland, ICPR January 1996.

The Alive system senses only a single user, and addresses only a constrained virtual world environment. Because the user is immersed in a virtual world, the context for the interaction is straight-forward, and, simple vision and graphics techniques can be employed. Sensing multiple users in an unconstrained real-world environment, and providing behavior-driven output in the context of that environment presents more complex

vision and graphics problems stemming from the requirement of real world interaction that are not addressed in prior art systems.

The Alive system fits a specific geometric shape model, such as a Gaussian ellipse, to a description representing the human user. The human shape model is referred to as a "blob." This method of describing shapes is generally inflexible. The Alive system uses a Gaussian color model which limits the description of the users to one dominant color. Such a limited color model limits the ability of the system to distinguish among multiple users.

The prior art system, supra, by Azarbayejani uses a self-calibrating blob stereo approach based on a Gaussian color blob model. This system has all of the disadvantages of inflexibility of the Gaussian model. The self-calibrating aspect of this system may be applicable to a desktop setting, where a single user can tolerate the delay associated with self-calibration. In a kiosk setting, it would be preferable to calibrate the system in advance so it will function immediately for each new user.

The prior art systems use the placement of the user's feet on the ground plane to determine the position of the user within the interaction space. This is a reasonable approach in a constrained virtual-reality environment, but this simplistic method is not acceptable in a real-world kiosk setting where the user's feet may not be visible due to occlusion by nearer objects in the environment. Furthermore, the requirement to detect the ground plane may not be convenient in practice because it tends to put strong constraints on the environment.

It remains desirable to have an interface paradigm for a computerized kiosk in which computer vision techniques are used not only to sense users but also to interact with them.

### SUMMARY OF THE INVENTION

The problems of the public user interface for computers are solved by the present invention of a computer vision technique for the visual sensing of humans, the modeling of response behaviors, and audiovisual feedback to the user in the context of a computerized kiosk.

The invention, in its broad form, resides in a computerized method and apparatus for interacting with a moving object in a scene observable with a camera, as recited in claims 1 and 10 respectively.

In a preferred embodiment described hereinafter, the kiosk has three basic functional components: a visual sensing component, a behavior module and a graphical/audio module. It has an optional component that contains three dimensional information of the environment, or observed scene. These components interact with each other to produce the effect of a semi-intelligent reaction to user behavior. The present invention is implemented using real-time visual sensing (motion detection, color tracking, and stereo ranging),

and a behavior-based module to generate output depending on the visual input data.

## BRIEF DESCRIPTION OF THE DRAWINGS

A more detailed understanding of the invention may be had from the following description of a preferred embodiment, given by way of example, and to be understood with reference to the accompanying drawing wherein:

- ♦ FIG. 1 is a block diagram of a public computerized user interface;
- ♦ FIG. 2 shows a kiosk and interaction spaces;
- ♦ FIG. 3 is shows a block diagram of the kiosk;
- ♦ FIG. 4 shows a four zone interaction space;
- ♦ FIG. 5 shows a flow diagram of an activity detection program;
- ♦ FIG. 6 is a block diagram of a behavior module process; and
- ♦ FIG. 7 shows an arrangement for stereo detection of users.

## DETAILED DESCRIPTION

Referring now to the figures, FIG. 1 shows a public computer user interface 10. The user interface 10 has a sensing module 15 which takes in information from a real world environment 20, including the presence and actions of users. The information is processed in a behavior module 25 that uses a three dimensional model 30 to determine proper output through a feedback module 35. The three dimensional model 30 of a real world environment 20, also referred to as a scene, includes both metric information and texture that reflect the appearance of the world.

FIG. 2 shows a kiosk 50 with a display screen 55 for the users of the kiosk, and a plurality of cameras 60, 65, 70 which allow the kiosk 50 to detect the presence of the users. Three cameras are shown, but a single camera, or any multiple of cameras may be used. A first camera 60 is aimed at an area on the floor. The "viewing cone" of the first camera 60 is defined to be a first interaction space 75. Second and third cameras 65, 70 are aimed to cover a distance out into the kiosk environment. In the present embodiment of the invention the second and third cameras 65, 70 are aimed out to 50 feet from the kiosk. The space covered by the second and third cameras 65, 70 is a second interaction space 80.

The kiosk 50 includes a visual sensing module 15 which uses a number of computer vision techniques, activity detection, color recognition, and stereo processing, to detect the presence or absence, and the posture of users in the interaction spaces 75, 80. Posture includes attributes such as movement and three dimensional spatial location of a user in the interaction spaces 75, 80. The kiosk digitizes color frames from the cam-

eras that are used by the visual sensing module 15 in the kiosk.

FIG. 3 is a block diagram of the kiosk 50. The kiosk 50 has input devices which include a plurality of cameras 100 coupled to digitizers 105 and output devices which may, for example, include a speaker 110 for audio output and a display screen 115 for visual output. The kiosk 50 includes a memory/processor 120, a visual sensing module 15, a behavior module 25, and a feedback module 35. The kiosk may also include a three dimensional model 30 representative of the scene 20. The visual sensing module 15 includes a detection module 125, a tracking module 130, and a stereo module 135 components which will be more fully described below.

The activity detection module 125 which uses computer vision techniques to detect the presence and movement of users in the interaction spaces of Figure 2. The kiosk 50 accepts video input of the interaction spaces from one or more cameras. In the first embodiment of the invention, the activity detection module 125 accepts video input from a single camera 60 which is mounted so that it points at the floor, as shown in FIG. 2. In operation, the activity detection module 125 examines each frame of the video signal in real-time to determine whether there is a user in the first interaction space 75, and if so, the speed and direction with which the person is moving. The activity detection module sends a message, or notification, to the behavior module every time a moving object enters and exits the first interaction space 75.

The first interaction space 75 is partitioned into one or four zones in which "blobs" are independently tracked. Where a regular camera lens is used, one zone is appropriate. Where a wide-angle or fisheye lens is used, four zones, as shown in FIG. 4, are used. The four zones are defined as a center zone 250, a left zone 255, a right zone 260, and a back zone 265. In the four zone mode, computations for activity detection are performed independently in each zone. The extra computations make the activity detection program more complex but allow more accurate estimation of the velocity at which the user is moving.

When there are four zones in the first interaction space 75, the kiosk is primarily concerned with blobs in the center zone 250, i.e. potential kiosk users. When a blob first appears in the center zone 250, the blob in a peripheral zone from which the center blob is most likely to have originated is selected. The velocity of this source blob is assigned to the center blob. The activity detection program applies standard rules to determine which peripheral zone (Right, Left or Back) is the source of the blob in the center zone 250.

The activity detection module compares frames by finding the difference in intensity of each pixel on the reference frame with the corresponding pixel in a new digitized frame. Corresponding pixels are considered to be "different" if their gray levels differ by more than a first

pre-defined threshold.

The activity detection program distinguishes between a person and an inanimate object, such as a piece of litter, in the first interaction space 75 by looking for movement of the object's blob between successive images. If there is sufficient movement of the object's blob between successive frames, the object is assumed to be animate. There is "sufficient motion" when the number of pixels that differ in successive images is greater than a second threshold.

FIG. 5 shows a flow chart of the operation of the activity detection program. At initialization of the activity detection program, block 400, the first interaction space 75 is empty and the kiosk 50 records a frame of the floor in the first interaction space 75. This initial frame becomes the reference frame 455 for the activity detection program. Approximately every 30 milliseconds, a new frame is digitized, block 400. A comparison, block 405, is then made between this new frame and the reference frame 455. If the new frame is sufficiently different from the reference frame 455 according to the first predefined pixel threshold value, the activity detection module presumes there is a user in the first interaction space 75, block 410. If the new frame is not sufficiently different, the activity detection program presumes that no one is in the first interaction space 75, block 410. If the activity detection program determines that there is a user in the first interaction space 75, the activity detection program sends a message to the behavior module 25, block 420. If the activity detection program determines that there is no person in the first interaction space 75, the behavior module is sent a notification, block 415, and a new frame is digitized, block 400.

If at block 410, the difference is greater than the first predefined threshold, a notification is also provided to the behavior module, block 420. The message indicates that something animate is present in the interaction space 75. At the same time, a frame history log 425 is initialized with five new identical frames which can be the initial frame (of block 400), block 430. A new frame, captured between significant intervals (approximately once every 10 seconds in the present embodiment), block 435, is then compared with each frame in the log to determine if there is a difference above a second threshold, block 440. The second threshold results in a more sensitive reading than the first threshold. If there is a difference above the second threshold, block 445, the frame is added to the frame history, block 430, a five frame-rotating buffer. The steps of blocks 430, 440, and 445 then repeat which indicates that an animate object has arrived. If there is a difference below the second threshold, block 445, the frame is blended with the reference frame, block 450, to create the new reference frame 455. The end result of the activity detection program is that the background can be slowly evolved to capture inanimate objects that may stray into the environment, as well as accommodate slowly changing characteristics such as lighting changes.

If there is a moving object in the first interaction space 75, the activity detection program computes the velocity of that object by tracking; in each video frame, the location of a representative point of the object's blob, or form. The blob position in successive frames is smoothed to attenuate the effects of noise using known techniques such as Kalman filtering. The activity detection program maintains a record of the existence of potential users in the kiosk interaction space 75 based on detected blobs.

### Velocity Computation

The activity detection program computes the velocity of users moving in the first interaction space 75 by tracking blob positions in successive frames. Velocity is used to indicate the "intent" of the blob in the first interaction space 75. That is, the velocity is used to determine whether the blob represents a potential user of the kiosk.

Velocity is computed as a change in position of a blob over time. For the velocity calculation, the blob position is defined as the coordinates of a representative point on the leading edge of the moving blob. When there is only one zone in the interaction space, the representative point is the center of the front edge of the blob. When there are four zones in the interaction space, the representative point differs in each zone. In the center and back zones, the point is the center of the front edge of the blob 252, 267. In the left zone, the point is the front of the right edge of the blob 262. In the right zone, the point is the front of the left edge of the blob 257. The velocities of blobs are analyzed independently in each zone.

### Behavior module

The behavior module 25, shown in FIG. 6, uses the output of the visual module 15 as well as a priori information such as the three dimensional model of the environment 30 to formulate actions. The behavior module 25 uses a set of rules (with the potential for learning from examples) as a means of reacting to user behavior in a manner that can be perceived as being intelligent and engaging. The mechanism for reacting to external visual stimuli is equivalent to transitioning between different states in a finite state machine based on known (or learnt) transition rules and the input state. As a simple example, the behavior module 25 can use the output of the detection module 125 to signal the feedback module 35 to acknowledge the presence of the user. It can take the form of a real time talking head in the display screen 55 saying "Hello." Such a talking head is described in "An Automatic Lip-Synchronization Algorithm for Synthetic Faces," Keith Waters and Tom Levergood, Proceedings of the Multimedia ACM Conference, September 1994, pp. 149 - 156. In a more complicated example, using the output of the stereo module 135

(which yields the current three dimensional location of the user/s), the behavior module 25 can command the talking head to focus attention on a specific user by rotating the head to fixate on the user. In the case of multiple users, the behavior module 25 can command the talking head to divide its attention amongst these users. Heuristics may be applied to make the kiosk pay more attention to one user than the other (for example, based on proximity or level of visual activity). In another example, by using both the stereo module 135 and three dimensional world information 30, the behavior module 25 can generate directional information, either visually or orally, to the user based on the user's current three dimensional location.

### Color Blob

Color blobs are used to track the kiosk users as they move about the interaction space. The distribution of color in a user's clothing is modeled as a histogram in the YUV color space. A color histogram detection algorithm used by the present invention is described in the context of object detection in "Color Indexing" by Michael J. Swain and Dana H. Ballard, International Journal of Computer Vision, 7:1, 1991, pp. 11 - 32. In the present invention, the color histogram method is used for user tracking and is extended to stereo localization.

Given a histogram model, a histogram intersection algorithm is used to match the model to an input frame. A back projection stage of the algorithm labels each pixel that is consistent with the histogram model. Groups of labeled pixels form color blobs. A bounding box and a center point are computed for each blob. The bounding box and the center point correspond to the location of the user in the image. The bounding box is an x and y minimum and maximum boundary of the blob. The color blob model has advantages for user tracking in a kiosk environment. The primary benefit is that multiple users can be tracked simultaneously, as long as the users are wearing visually distinct clothing. The histogram model can describe clothing with more than one dominant color, making it a better choice than a single color model. Histogram matching can be done very quickly even for an NTSC resolution image (640 by 480 pixels), whereby a single user may be tracked at 30 frames per second. Color blobs are also insensitive to environmental effects. Color blobs can be detected under a wide range of scales, as the distance between the user and the camera varies. Color blobs are also insensitive to rotation and partial occlusion. By normalizing the intensity in the color space, robustness to lighting variations can be achieved. The center locations, however, of detected color blobs are significantly affected by lighting variation. Use of color for tracking requires a reference image from which the histogram model can be built. In the architecture of the present embodiment of the invention, initial blob detection is

provided by the activity detection module, which detects moving objects in the frame. The activity detection module assumes that detected blobs correspond to upright moving people, and samples pixels from the central region of the detected blob to build the color histogram model.

### Stereo

Through stereo techniques, true three dimensional information about user location can be computed from cameras in an arbitrary position relative to the scene. Stereo techniques require frames from two or more cameras be acquired concurrently, as shown in FIG. 7. This is a known method for computing detailed descriptions of scene geometry. In a classical approach, frames acquired from two cameras are processed and the correspondences between pixels in the pair of frames are determined. Triangulation is used to compute the distance to points in the scene given correspondences and the relative positions of the cameras. In the classical approach, a high level of detail retires excessive computational resources. The method of the present embodiment is based on a simpler, object-based version of the classical stereo technique. Moving objects are tracked independently using color or motion blobs in images obtained from synchronized cameras. Triangulation on the locations of the moving objects in separate views is used to locate the subjects in the scene. Because tracking occurs before triangulation, both the communication and computational costs of dense stereo fusion are avoided.

The triangulation process is illustrated in Figure 7. Given the position of a blob 700 in a first camera image 702, the position of the user 705 is constrained to lie along a ray 710 which emanates from a first camera 715 through the center of the blob 700 and into the scene. Given the position of a second blob 712 in a second camera image 720, the position of the user 705 is constrained to lie along a second ray 725. The user 705 is located at the intersection of the first ray 710 and the second ray 725 in the scene. In actual operation, noise in the positions of the blobs 700, 712 makes it unlikely that the two rays 710, 725 will intersect exactly. The point in the scene where the two rays 710, 725 are closest is therefore chosen as the three dimensional location of the user 705.

In a preferred embodiment of the kiosk system, a pair of verged cameras with a six foot baseline, i.e. separation between the cameras, is used. The stereo approach depends on having calibrated cameras for which both the internal camera parameters and relationship between camera coordinate systems are known. A standard non-linear least squares algorithm along with a calibration pattern to determine these parameters off-line are used.

Camera synchronization is achieved by ganging the external synchronization inputs of the cameras

together. Barrier synchronization is used to ensure that the blob tracking modules that process the camera images begin operation at the same time. Synchronization errors can have a significant effect on conventional stereo systems, but blobs with large size and extent make stereo systems much more robust to these errors. 5

It is to be understood that the above-described embodiments are simply illustrative of the principles of the invention. The present invention has been described in the context of a kiosk however alternative embodiments could be automated teller machines (ATMs), advanced multimedia TV, or office desk computers. Various and other modifications and changes may be made by those skilled in the art which will embody the principles of the invention and fall within the scope thereof. 10 15

### Claims

1. A computerized method for interacting with a moving object or person in a scene observable with a camera, comprising the steps of: 20

determining a posture of the moving object by comparing successive frames of the scene;  
outputting information which can be sensed by the moving object depending on the posture of the object as determined from the comparison of the successive frames. 25

2. The method of claim 1, wherein the posture of the moving object includes a position of the moving object, wherein further the position is determined in three dimensional space, and multiple cameras are used to observe the scene. 30

3. The method of claim 1, wherein the scene includes a plurality of moving objects, the method including observing dominant colors of the plurality of moving objects to interact independently with any of the moving objects. 35 40

4. The method of claim 1, wherein the outputted information includes audible and visible signals, further comprising: 45

displaying a talking head on a display terminal, the method including controlling the orientation of the talking head depending on the posture of the moving object.

5. The method of claim 4, wherein the step of synchronizing audible signals with the orientation of the talking head and dependent on the posture of the moving object. 50

6. The method of claim 1 further comprising:  
repeatedly storing a previous frame of the 55

scene in a buffer if a difference between the previous frame and a next frame is greater than a predetermined value;

determining the posture of the moving object by analyzing the frames stored in the buffer.

7. A computerized apparatus for interacting with a moving object or person in a scene observable with a camera, comprising:

means for determining a posture of the moving object by comparing successive frames of the scene;

means for outputting information which can be sensed by the moving object depending on the posture of the object as determined from the comparison of the successive frames.

8. A computerized interface for interacting with people, comprising:

a camera measuring a region of an arbitrary physical environment as a sequence of images; and

means for detecting a person in the region from the sequence of images to identify the person as a target for interaction.

9. The interface of Claim 8, further comprising:

means for rendering audio and visual information directed at the detected person, further comprising:

means for determining a velocity of the person in the region; and wherein a content of the rendered audio and video information depends on the velocity of the person, wherein the means for rendering includes a display system displaying an image of a head including eyes and a mouth with lips, the display system directing an orientation of the head and a gaze of the eyes at the detected person while rendering the audio information synchronized to movement of the lips so that the head appears to look at and talk to the person.

- 50 10. The interface of Claim 9, further comprising:

means for determining a position and an orientation of the person in the region relative to a position of the camera, further comprising:

means for rendering audio and video information directed at the detected person, a content of the rendered information

depending upon the determined position and the determined orientation of the person in the region.

of the cameras.

11. The interface of Claim 8, further comprising:

a memory, coupled to the means for detecting, storing data representing a three-dimensional model of the physical environment for determining a position of the person in the region relative to objects represented in the three-dimensional model, further comprising:

means for rendering audio and video information, a content of the rendered information depending upon the determined position of the person.

12. The interface of Claim 8, wherein the sequence of images includes a reference image and a target image, each image being defined by pixels, the pixels of the reference image having a one-to-one correspondence to the pixels of the target image; and further comprising:

means for comparing the reference image to the target image to identify a group of adjacent pixels in the reference image that are different from the corresponding pixels in the target image, the identified group of pixels representing the person, wherein the means for comparing compares an intensity of each pixel of the reference image to an intensity of each corresponding pixel in the target image, and the means for detecting detects the presence of the person in the region when the intensities of at least a pre-defined number of the pixels of the reference image differ from the intensities of the corresponding pixels of the target image.

13. The interface of Claim 12 further comprising:

means for blending the target image with the reference image to generate a new reference image when less than a pre-defined number of the pixels of the reference image differ from the corresponding pixels of the target image.

14. The interface of Claim 8 further comprising:

a second camera spaced apart from the other camera, the second camera measuring the region as a second sequence of images, further comprising:

means for determining an approximate three-dimensional position of the person in the region from the sequences of images

15. The interface of Claim 8 further comprising:

means for rendering audio and visual information, the rendered audio and video information interacting in turn with a plurality of detected persons.

16. The interface of Claim 8, wherein the sequence of images includes a reference image and a target image, each image being defined by pixels, the pixels of the reference image having a one-to-one correspondence to the pixels of the target image; and further comprising:

means for comparing the reference image to the target image to identify a plurality of groups of adjacent pixels in the reference image that are different from the corresponding pixels in the target image, each identified group of pixels representing one of a plurality of detected persons.

17. The interface of Claim 16 further comprising:

means for determining a distribution of colors in each of the group of pixels, each color distribution uniquely identifying one of the plurality of persons, further comprising:

means for concurrently tracking movements of each person independently in the region by the color distribution that uniquely identifies that person.

18. A computerized interface for interacting with people, comprising:

a camera measuring a region of an arbitrary physical environment as a sequence of images; and

means for rendering audio and video information directed at a person detected in the region from the sequence of images to interact with the person.




FIG. 1






FIG. 2




FIG. 3




FIG. 4

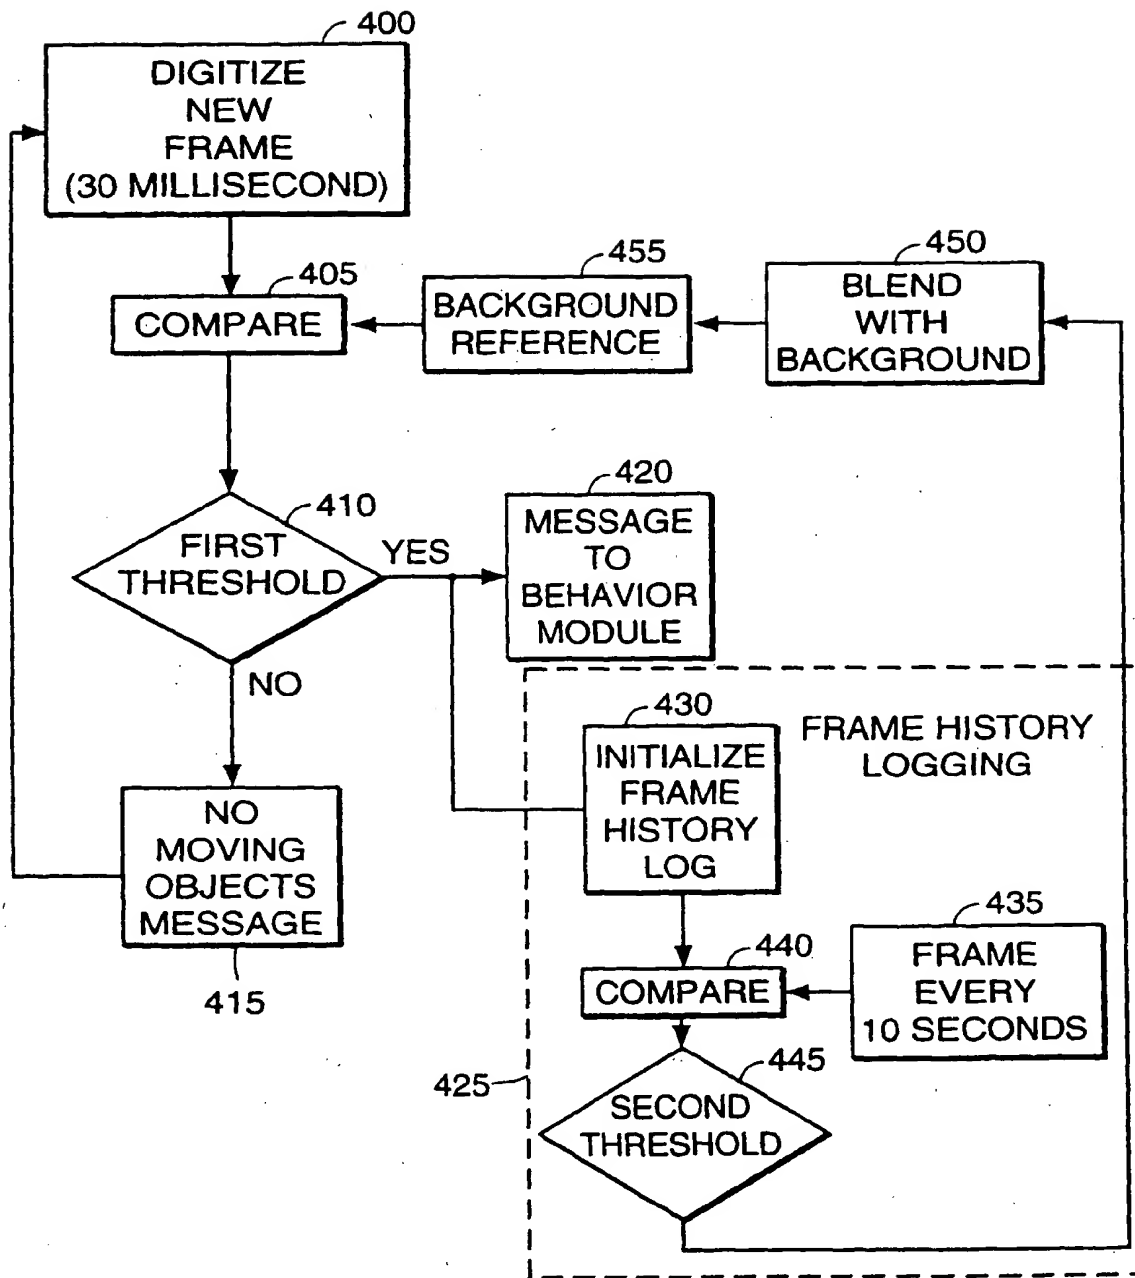


FIG. 5




FIG. 6




FIG. 7



European Patent  
Office

# EUROPEAN SEARCH REPORT

Application Number  
EP 98 10 6380

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
X	EP 0 582 989 A (ISTITUTO TRENINO DI CULTURA) 16 February 1994 * page 3, line 34 - page 4, line 51 * * page 5, line 15-32 * * figure 1 *	1,7,8,18	G06K11/08
Y	---	2,4,5,14	
A	---	6,10-12	
X	US 5 247 433 A (KITAURA WATARU ET AL) 21 September 1993 * column 13, line 36 - column 14, line 30 * * figure 20 *	1,7,8,18	
A	---	11	
Y	PATENT ABSTRACTS OF JAPAN vol. 097, no. 008, 29 August 1997 & JP 09 097337 A (FUJI HEAVY IND LTD), 8 April 1997 * abstract *	2,14	
A	---	1,6-8,12,18	TECHNICAL FIELDS SEARCHED (Int.Cl.6)
Y	EP 0 694 833 A (AT & T CORP) 31 January 1996 * page 3, line 45-55 * * page 4, line 11-29 * * page 7, line 37-42 * * figures 1-3 *	4,5	G06K G06F
A	---	9,10,15	
A	"METHOD FOR EXTRACTING FACIAL FEATURES BY USING COLOR INFORMATION" IBM TECHNICAL DISCLOSURE BULLETIN, vol. 38, no. 10, 1 October 1995, pages 163-165, XP000540455 * the whole document *	3,17	
The present search report has been drawn up for all claims			
Place of search <b>THE HAGUE</b>		Date of completion of the search <b>31 July 1998</b>	Examiner <b>Baldan, M</b>
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons * : member of the same patent family, corresponding document	

EPO FORM 1503 03/82 (P04C01)



European Patent  
Office

# EUROPEAN SEARCH REPORT

Application Number  
EP 98 10 6380

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
X,P	EP 0 818 726 A (SIEMENS AG) 14 January 1998 * column 3, line 52 - column 5, line 48 * * column 6, line 5-47 * * figures 1,2,3B-4B *	1,7,8,18	
A	-----	10,11	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
The present search report has been drawn up for all claims			
Place of search <b>THE HAGUE</b>		Date of completion of the search <b>31 July 1998</b>	Examiner <b>Baldan, M</b>
<p><b>CATEGORY OF CITED DOCUMENTS</b></p> <p>X : particularly relevant if taken alone  Y : particularly relevant if combined with another document of the same category  A : technological background  O : non-written disclosure  P : intermediate document</p> <p>T : theory or principle underlying the invention  E : earlier patent document, but published on, or after the filing date  D : document cited in the application  L : document cited for other reasons  &amp; : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03.92 (P04C01)